

**United States v. Shacar
21-cr-30028-MGM
EXHIBIT “P”**

UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS

UNITED STATES OF AMERICA)
)
)
v.) Docket No. 20-cr-10012-IT
)
)
)
PAUL BATEMAN)

MOTION TO SUPPRESS

Exhibit H:

Declaration of Professor Steven Murdoch

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA
Alexandria Division

UNITED STATES OF AMERICA,
Plaintiff,
v.
ZACKARY ELLIS SANDERS,
Defendant

Case No.: 1:20-cr-00143

DECLARATION OF PROFESSOR STEVEN MURDOCH

I, Professor Steven Murdoch, Declare under penalty of perjury that:

1. I am currently Professor of Security Engineering at University College London. My research is focused on information security, particularly Internet privacy and payment systems security.
2. I have a Ph.D. from the Security Group of the University of Cambridge Department of Computer Science and Technology. My doctoral research, completed in 2007, focused on the traffic analysis of anonymous communications systems, in particular Tor. I have worked with the developers of Tor since 2004, and I created the first version of the Tor Browser software in 2008. I continue to work with the Tor Project, the 501(c)(3) non-profit organization responsible for the development of the Tor network and associated software. I helped the Tor Project assess and improve the security and usability of Tor.

3. I am a member of Christ's College Cambridge and a Royal Society University Research Fellow. I am a Fellow of the Institution of Engineering and Technology (IET) and the British Computer Society (BCS).
4. I have published a number of peer-reviewed research papers in the fields of anonymous communications and privacy-enhancing technologies, including Tor.
5. My CV is attached as Exhibit A.

Scope of declaration

6. This declaration will discuss how the Tor network operated and how activity on the Tor network could be traced as of May 23, 2019 – the date of the alleged visit in this case. I have restricted this declaration to the use of websites running as Tor Onion Services (also known as Tor Hidden Services), where the website is accessible only through the Tor network. I also will only discuss version 2 Tor Onion Services, in which the Onion Service address is 16 characters followed by “.onion” as the website in question for this case is a version 2 Onion Service (because according to a screenshot disclosed in discovery, its address had 16 characters before the “.onion,” not 56).
7. The vast majority of usage of the Tor network (about 98% as of June 2021) is to visit websites and other Internet services which are also accessible directly over the Internet. I will not discuss this use of Tor in my declaration because it is not the type of usage alleged in this case. Similarly, I will not discuss version 3 Tor Onion Services, in which the Onion Service address is 56 characters followed by “.onion.”

Operation of Tor Onion Services

8. When a user wishes to visit a Tor Onion Service, various computers (nodes) within the Tor network are involved. Some of these nodes are selected at random by the Tor software running on the user's computer, and the others are selected at random by the Tor software running on the Onion Service.
9. The Tor software on the user's computer will select at random a node to act as the rendezvous node and send the address of the rendezvous node to the Onion Service.

To protect the anonymity of the Onion Service (i.e., its IP address),¹ its connection is made through a three-node Tor circuit. Similarly, to protect the anonymity of the user, its connection to the rendezvous node is through a three-node Tor circuit.

10. Next, the Onion Service will connect to the rendezvous node through a three-node Tor circuit (excluding the rendezvous node). The rendezvous node is responsible for connecting together the two circuits: one created by the user’s computer and one by the Onion Service. The Onion Service and user can then exchange information, i.e., the user’s web browser can request to view a web page and other data, and the Onion Service can respond with the requested content, without disclosing each other’s IP addresses.
11. When the Onion Service and user’s computer are connected, there are six nodes between them: the guard node selected by the user’s computer, one middle node, the rendezvous node, two middle nodes, and the guard node selected by the Onion Service. Only the Onion Service’s guard node is aware of the Onion Service’s IP address, and only the user’s guard node is aware of the user’s IP address.

Improvements in the Tor network’s design and operation prior to May 2019 have made traffic analysis less reliable and more difficult to execute

12. Traffic analysis of Tor has always been unreliable because of security features which were present in Tor since its creation. For example, data sent between nodes has always been padded so that everyone’s communications are the same size. As a result of this and other security measures, in 2012, NSA documents showed that “with manual analysis we can de-anonymize a very small fraction of Tor users.”²
13. The Tor Project is continually improving the Tor network and associated software to enhance the level of privacy and usability offered to Tor users. This work is

¹ The IP address could allow the physical location of the Onion Service to be found.

² Tor Stinks, presentation by National Security Agency, dated June 2012. Published by The Guardian, October 4, 2013. Available at <https://edwardsnowden.com/2013/10/04/tor-stinks-presentation/> (emphasis in original)

conducted in collaboration with academic researchers and other organizations. To help build trust in the development process, discussions of proposed and implemented changes to the Tor network are conducted publicly through academic literature, mailing lists, Internet chat, and websites maintained by the Tor Project. As such, the changes discussed here and the time frame in which they occurred are widely known.

14. Since September 2014, the Tor software will select one guard node and keep this selection for as long as possible to minimize the number of computers with knowledge of IP addresses. For this reason, if a Tor node collected IP addresses of users connecting to it, only a very small proportion of Tor users would be identified, no matter how long this node was run. It takes 68 days before a new node can run at full capacity, which is a measure put in place to prevent someone from quickly creating new nodes in an attempt to collect more IP addresses.
15. The likelihood that the Tor software will select any given node is proportional to that node's contribution to the capacity of the Tor network. Therefore, as the capacity of the Tor network has grown, the less likely it is that any one node will be selected, so it makes it less likely that a particular node can be used to observe users of the Tor network. For example, in 2012, NSA documents showed that "with manual analysis we can de-anonymize a very small fraction of Tor users."³ Since then, the network capacity has grown from about 10 gigabits per second to 200 gigabits per second, which has further increased the difficulty of performing traffic analysis by a factor of 20.
16. Tor's encryption ensures that the content of data sent to and from a user's guard node is impossible to match to the corresponding content of data sent to or from the Onion Service's guard node. In other words, looking at the content of data the user sends to their guard node will be of no help in identifying with which Onion Service that user is communicating.

³ Tor Stinks, presentation by National Security Agency, dated June 2012. Published by The Guardian, October 4, 2013. Available at <https://edwardsnowden.com/2013/10/04/tor-stinks-presentation/> (emphasis in original)

17. Traffic analysis is a technique to attempt nevertheless to identify which user is communicating with which Onion Service by comparing patterns of when and how much data is sent (as opposed to looking at the content of the data, which is not visible to observers).
18. Prior to 2016, traffic analysis on Tor was unreliable, but there were concerns that it might be possible in some cases because the rate at which data was sent between the user and their guard node was similar to the rate at which data was sent between the Onion Service and its guard node. Prior to 2016, if someone were able to observe those two links simultaneously, patterns in the two rates of data transfer may be sufficiently similar to identify the link with some degree of confidence. That was no longer the case beginning sometime in 2016.
19. In 2016 Tor introduced an extension to its padding feature, to obscure patterns in the rate of data transfer to and from guard nodes. The Tor software randomly sends additional padding data to and from guard nodes which does not continue on through the rest of the Tor circuit. In so doing, the Tor software hides patterns in data rates that might otherwise allow traffic analysis to link users to the Onion Services they are visiting. Padding will cause traffic analysis to introduce more errors, both false-positives (where a user is incorrectly identified as having visited the Onion Service) and false-negatives (where a user is incorrectly identified as not having visited the Onion Service). In an undated presentation published in 2014,⁴ GCHQ noted a high level of errors in performing traffic-analysis on Tor, even before the extended padding feature was introduced in 2016.
20. The nodes that make up the Tor network are operated by thousands of volunteers around the world, many of whom do not disclose their identities. In some cases, operators of Tor nodes were found to tamper with or observe traffic on the nodes. To prevent operators from being able to tamper with or observe traffic passing through

⁴ A potential technique to deanonymise users of the TOR network (undated). Published by Der Spiegel, December 28, 2014. Available at <https://edwardsnowden.com/2015/01/07/a-potential-technique-to-deanonymise-users-of-the-tor-network/>

their node, since 2014, the Tor Project has pro-actively identified Tor nodes that show suspicious behavior and then exclude them from the network. This has increased the difficulty of covertly performing traffic analysis on the Tor network since 2014.

21. Thus, prior to May 2019, numerous measures were implemented in the Tor software to make it even more difficult to use traffic-analysis to de-anonymize Tor users.

Law-enforcement likely had control over the Onion Service when the visit was recorded

22. Law enforcement can attempt to identify the IP address of a Tor user visiting an Onion Service in four scenarios:

- 1) Law enforcement controls neither the user's guard nor the Onion Service
- 2) Law enforcement can observe the user and observe the Onion Service
- 3) Law enforcement controls the Onion Service and some guard nodes
- 4) Law enforcement controls the Onion Service

23. Although I describe four scenarios, there are only two techniques for identifying the IP address of a user using Tor Browser properly: traffic-analysis (which can generate errors) or a Network Investigative Technique (which interferes with a user's computer).

Scenario 1: Law enforcement controls neither the user's guard nor the Onion Service

24. In the first scenario, law-enforcement can only use traffic analysis if both the Onion Service and the user happen to randomly select a guard node controlled or observed by law-enforcement. Because guard nodes are selected at random, the chance of simultaneously controlling or observing both guard nodes is very small. It is for this reason, I believe the NSA concluded in 2012 that reliably de-anonymizing Tor users was not feasible, and since then, the difficulty has increased.

Scenario 2: Law enforcement can observe the user and observe the Onion Service

25. In the second scenario, law-enforcement could use traffic analysis to confirm if a particular user was visiting an Onion Service. For example, in the complaint dated October 29, 2014, alleging that Blake Benthall operated the Silk Road 2.0 Onion

Service,⁵ the FBI noted similarities between Tor traffic observed through pen-register data at the defendant's residence and activity on the Onion Service in question. There are still likely to be errors in traffic-analysis because of Tor's padding obscuring patterns in the rates of data, leading to both false-positives and false-negatives.

Scenario 3: Law enforcement observes the Onion Service and some guard nodes

26. In the third scenario, law-enforcement can use traffic analysis to try to identify random visitors to an Onion Service (as opposed to specific targets). Traffic analysis is more likely to succeed than in the first scenario, but the likelihood of doing so is still very small. Firstly, it would only be possible to trace users who happen to select a guard under observation by law-enforcement. Secondly, even in the small number of cases where this has occurred, Tor's padding will cause traffic analysis performed through observing traffic to have a significant number of errors. Even if there are similarities in traffic patterns at a given guard node under observation and at the Onion Service, it could be that a different guard node is the correct match but is not under observation. As the number of visitors to an Onion Service increases, so will the likelihood of traffic analysis errors. Increasing the traffic to an Onion Service to frustrate traffic analysis is known as adding "cover traffic." To reduce the chances of errors, law-enforcement should wait until traffic-analysis indicates that the same user has visited an Onion Service multiple times.

Scenario 4: Law enforcement controls the Onion Service

27. In the fourth scenario, law-enforcement could use a Network Investigative Technique (NIT) to identify the IP address of a user visiting an Onion Service by forcing the user's computer to disclose its IP address by connecting directly to a law-enforcement server without using the Tor network.

⁵ United States of America v. Blake Benthall, Southern District of New York. October 29, 2014. Available at <https://www.justice.gov/sites/default/files/usao-sdny/legacy/2015/03/25/Benthall,%20Blake%20Complaint.pdf>

28. In order to use an NIT, law-enforcement must control the Onion Service prior to deploying the NIT. If law-enforcement controls the Onion Service, applying an NIT would be feasible and avoid the error-prone nature of traffic analysis.
29. An NIT used in such a scenario could be based around malware that executes on the user's computer and forces it to directly connect to the Internet (as opposed to connecting through Tor). The FBI used this technique to identify visitors to the Freedom Hosting servers.⁶ Such an NIT causes the user's computer to do something it would not otherwise do and therefore interferes with the user's computer.
30. Alternatively, the NIT may not involve malware code and instead force the user's Tor Browser to malfunction and connect directly through the Internet rather than sending network traffic through the Tor network. Such an NIT causes the Tor Browser on the user's computer to malfunction and reveal information that it was designed not to, and therefore interferes with the user's computer.
31. In summary, I do not believe the first scenario is plausible for this case. Traffic analysis is extremely unlikely to yield the hundreds of IP addresses submitted by the [REDACTED], nor give the [REDACTED] confidence that these IP addresses visited the Onion Service in question.
32. I, therefore, conclude that law-enforcement almost certainly controlled the Onion Service prior to May 23, 2019, and either used traffic analysis or an NIT to identify visitors to the Onion Service. As discussed above, a single identification using traffic analysis could very well be a false-positive error. An NIT necessarily interferes with a user's computer wherever it is located.

⁶ Feds Are Suspects in New Malware That Attacks Tor Anonymity, Kevin Poulsen, Wired News, August 5, 2013. Available at <https://www.wired.com/2013/08/freedom-hosting/>

Neither of the two papers that Special Agent Ford cites could explain what has happened in this case

33. For reasons discussed below, neither of the two papers cited by Special Agent Ford, in paragraph 7 of the declaration dated August 10, 2020, could explain what has happened in this case.

The existence of research on how to defend against a global passive adversary does not imply a global passive adversary is a realistic threat

34. Neither paper shows that a global-passive adversary is possible.

35. Information security research commonly makes use of unrealistic and hypothetical scenarios to help ensure that systems are secure in all realistic scenarios. This approach is analogous to that taken for systems where a failure would be dangerous, such as elevators. These are designed to withstand a weight far exceeding what is realistic (or even possible) so as to create a safety margin when they are used in the real world.

36. The global-passive adversary is a similar type of theoretical assumption. A global-passive adversary is a hypothetical single organization that can observe every computer and every network connection in the world simultaneously and record all information. This is impossible, and even the most capable intelligence agencies such as NSA or GCHQ cannot achieve this goal.

37. The global-passive assumption is a helpful thought experiment for research because if a system is designed to be secure even in the face of a global-passive adversary, then it should be safe in any realistic situation. The use of the global-passive adversary assumption in research does not, however, imply that such an adversary could ever exist in the real world.

The method of traffic-analysis that Special Agent Ford cites was not a feasible method in this case

38. The 2008 paper cited by Special Agent Ford in paragraph 7 of his declaration, by Chakravarty et al., assumed that the IP address of the user visiting the Onion Service was already known to law-enforcement and that law-enforcement were able to

manipulate network equipment that is carrying that user's traffic only to confirm that identification. This confirmatory technique only worked when the law-enforcement agency was on the same continent as the user, and even then, it was unreliable. Since 2008, this type of confirmatory technique has become much more difficult for the reasons discussed above. The 2008 paper is not relevant to understanding what law-enforcement could have done in 2019.

The Tor Project's advice has been taken out of context

39. The Special Agent's citation in paragraph 5, to the Tor Project's advice to its users has been taken out of context. It continued:

"First, Tor protects the network communications. It separates where you are from where you are going on the Internet. What content and data you transmit over Tor is controlled by you. If you login to Google or Facebook via Tor, the local ISP or network provider doesn't know you are visiting Google or Facebook. Google and Facebook don't know where you are in the world. However, since you have logged into their sites, they know who you are. If you don't want to share information, you are in control.

Second, active content, such as Java, Javascript, Adobe Flash, Adobe Shockwave, QuickTime, RealAudio, ActiveX controls, and VBScript, are binary applications. These binary applications run as your user account with your permissions in your operating system. This means these applications can access anything that your user account can access. Some of these technologies, such as Java and Adobe Flash for instance, run in what is known as a virtual machine. This virtual machine may have the ability to ignore your configured proxy settings, and therefore bypass Tor and share information directly to other sites on the Internet. The virtual machine may be able to store data, such as cookies, completely separate from your browser or operating system data stores. Therefore, these technologies must be disabled in your browser to use Tor safely.

That's where Tor Browser comes in. We produce a web browser that is preconfigured to help you control the risks to your privacy and anonymity while browsing the Internet. Not only are the above technologies disabled to prevent identity leaks, Tor Browser also includes browser extensions like NoScript and

Torbutton, as well as patches to the Firefox source code. The full design of Tor Browser can be read here. In designing a safe, secure solution for browsing the web with Tor, we've discovered that configuring other browsers to use Tor is unsafe."

40. First, Tor does not protect users from disclosing identifying information about themselves (e.g., name, email address), but in this case, there is no allegation that the Internet user ever logged into any website or revealed their email address in a way that allowed law-enforcement to identify them.
41. Second, the Tor Project advises users to use Tor Browser when connecting to Tor because the Tor Browser disables technologies that could disclose identifying information.
42. The advice from the Tor Project that Special Agent Ford cites does not discuss traffic analysis or the global passive adversary at all.
43. This guidance from the Tor Project does not mean that the theoretical scenarios Special Agent Ford outlined in his declaration are likely or even possible. In fact, when highly capable intelligence agencies have tried to de-anonymize Tor users, they have rarely succeeded.
44. The discussion by Special Agent Ford on tracing exit nodes, in paragraph 6, is also not relevant to the case. Exit nodes are only used when Tor users visit a website outside of the Tor network. When visiting a Tor Onion Service, an exit node is not used because data never leaves the Tor network.

Using a Tor search engine to visit an Onion Service website is easy

45. Paragraph 27 of the search warrant affidavit implies that searching to discover the address of an Onion Service is difficult and requires the use of a directory site.
46. In fact, entering "tor search" into the Tor Browser address bar (as of June 2021) will offer 5 different search engines for Tor Onion Services, allowing the user to find Onion Services matching particular keywords easily. Of these, at least two were available in May 2019. There is no need to use an index to find a Tor Onion Service because search engines are easily available from Tor Browser with just a few clicks.
47. Compared to the open Internet, there are fewer Onion Services, and so it is easier for a user to visit all sites returned for a search. For example, a search on the Torch search

engine for “Department of Justice” finds 61 Onion Service websites, so it would be easy to visit them all. In contrast, Google finds 107 million websites, and so there is no feasible way to visit all of these.

Visiting an Onion Service from a search engine does not imply the visitor was aware of the content of the website

48. Search engines, including for Onion Services, build an index of websites and associate these with keywords. This index is built by the search engine visiting publicly accessible parts of the website. Material on the website that is only visible after logging in will not be part of this index. Based on the screenshot of the target website’s homepage, the site’s homepage did not indicate the nature of the content available to logged-in users. Therefore the target site may appear in searches for innocuous keywords. Furthermore, when the site appears in search results, it will likely not be possible for a user to identify that there is illegal content without clicking on the search result to visit the website. For example, someone interested in BDSM could search for this term in a Tor search engine, but the results would not make clear what content is on the websites (see exhibit B attached, a search for “BDSM” as of June 2021 returning 10 results, any of which would be easy to click on).

49. Sites also have an incentive to encourage visitors from search engines, so the operator of the site may apply techniques to cause the site to appear more highly ranked for a wider range of keywords. This incentive results from the fact that some Onion Services are supported by advertising, just as with the open Internet. Furthermore, Onion Services who wish to protect the privacy of their operator and users visiting them may wish to increase the traffic to the website in order to obfuscate attempts to perform traffic analysis or other de-anonymization techniques (i.e., to add cover traffic). Search results may be intentionally misleading to attract more users who might not necessarily be interested in the content the website contains.

Directories will not necessarily indicate the content of a website

50. Paragraph 27 of the search warrant affidavit also implies that directories of Onion Service addresses give accurate and clear indications of whether these services contain unlawful material.

51. This is not necessarily the case. Some directories are open to editing by anyone and are not moderated. A person who wishes to promote an Onion Service may list the address while not indicating that the content is unlawful. As with the discussion of search engines above, this could be because, for example, the person wants to attract more users who are not motivated by illegal content but may nevertheless visit.

The difficulty of discovering the IP address of an Onion Service has nothing to do with how easy it is to visit an Onion Service

52. Paragraph 14 of the search warrant affidavit could be interpreted as saying that it is very difficult for users to view and visit Onion Services because they cannot find the IP address of the Onion Service using public lookups. However, this interpretation would not be correct.

53. Tor allows users to visit an Onion Service without knowing that service's IP address. Visiting an Onion Service from the Tor Browser is as simple as clicking on a link.

54. On the open Internet, law-enforcement can use a website's address (e.g., www.justice.gov) to identify its IP address. Tor does not allow law enforcement to do this for Onion Services. As a result, it is difficult, including for the reasons discussed above, for law enforcement to identify and locate who is running an Onion Service. That has nothing to do with how easy it is for a user to visit a Tor Onion Service because Tor allows users to visit an Onion Service without knowing its IP address.

A recorded visit to an Onion Service does not imply there was intent to visit that Onion Service or view the content on it

55. Just because someone visits a website, it does not imply that they intended to visit that website or view the content on that website.

56. It is common that a single website will contain parts of other websites. For example, on visiting the CNN homepage, content is downloaded from 70 different domain names, mostly for the purposes of showing advertisements or collecting statistics about who is visiting the site. The CNN homepage includes code that causes the web browser to download and display an image from Google to form part of an advertisement.

57. Onion Services can operate similarly. For example, an Onion Service (A) could include code that instructs Tor Browser to download images or other content from a different Onion Service (B), which then may or may not be displayed. If Onion Service B was under surveillance by law-enforcement, whether through an NIT or traffic analysis, a user visiting Onion Service A would also be seen to be visiting Onion Service B, even though the user only meant to visit Onion Service A. In the implementation of an NIT or traffic-analysis that I am aware of, it would not be possible for law-enforcement to distinguish someone directly visiting Onion Service B from someone who actually visits Onion Service A, which then triggers the indirect visit to Onion Service B.

58. Someone may include part of one Onion Service website within another to inflate the number of visitors to the website, to create cover traffic, to advertise content, or to collect statistics on who is visiting the site. The act of someone's browser visiting an Onion Service is not an indication of an intent to visit that website.

DONE this day, June 21, 2021.



Professor Steven Murdoch
Cambridge, UK

Exhibit A

PROFESSOR STEVEN J. MURDOCH

Address: Computer Science Department
University College London
Gower Street, London, WC1E 6BT **Email:** s.murdoch@ucl.ac.uk
Homepage: <https://murdoch.is/>

Education

2002–2007	University of Cambridge, Computer Laboratory (UK) – PhD in Computer Science Thesis: <i>Covert Channel Vulnerabilities in Anonymity Systems</i>
1998–2002	University of Glasgow (UK) – BSc Honours in Software Engineering (1st Class)

Professional History

Oct 20-	Professor (proleptic appointment), Computer Science, University College London
Oct 16-Oct 20	Reader (proleptic appointment), Computer Science, University College London
Aug 14-	Principal Research Fellow, Computer Science, University College London
Nov 13-	Innovation Security Architect, OneSpan
Dec 12-Jul 14	Research Fellow, Computer Laboratory, University of Cambridge
Jan 09-Nov 12	Senior Research Associate, Computer Laboratory, University of Cambridge
Aug 07-Sep 13	Chief Security Architect, Cronto
Aug 07-Dec 08	Research Associate, Computer Laboratory, University of Cambridge
Aug 06-Jun 07	Research Assistant, Computer Laboratory, University of Cambridge

Expert Witness

- Expert witness for Worcester Police Force, 2007
- R v Patel, Croydon Crown Court, 2008
- Job v Halifax PLC, case number 7BQ00307, Nottingham County Court, 2009
- Kaae v HSBC, London Mercantile Court, 2011
- R v Fisher, Camberwell Green Youth Court, 2014
- R v Vincent, Winchester Crown Court, 2014
- Expert witness relating to an application for third-party disclosure, 2014
- Ongoing case related to attribution of Internet traffic, 2019

Other Appointments and Affiliations

Fellowships:

- Fellow of the Institution of Engineering and Technology – FIET (2016–)
- Fellow of the British Computer Society – FBCS (2016–)
- Bye-Fellow, Christ’s College Cambridge (2014–)
- Research Fellow, Christ’s College Cambridge (2008–2014)

Journal Editor:

- Proceedings on Privacy Enhancing Technologies (2015)
- IEEE Internet Computing (special edition in 2013)
- SpringerBriefs in Cybersecurity (2012–)

International Award Chair:

- Andreas Pfitzmann (Privacy Enhancing Technologies Symposium) Award (2019)

International Conference Sponsorship Chair: Privacy Enhancing Technologies Symposium (2016–)

International Conference Program Chair: Privacy Enhancing Technologies Symposium (2014–2015)

International Conference General Chair: Financial Cryptography (2011)

International Conference Programme Committee Member:

- ACM CHI Conference on Human Factors in Computing Systems (2019)
- IEEE European Symposium on Security and Privacy (2019)
- Financial Cryptography (2010, 2016, 2018)
- IFIP Summer School (2008, 2016, 2017)
- Network and Distributed System Security Symposium (2017)
- ACM Conference on Computer and Communications Security (2007, 2008, 2010, 2011, 2016)
- Annual Privacy Forum (2014)
- USENIX Free and Open Communications on the Internet (2013)
- USENIX Security (2012)
- European Symposium on Research in Computer Security (2011)
- Privacy Enhancing Technologies Symposium (2007, 2008, 2009, 2011)
- Workshop on Foundations of Security and Privacy (2010)
- Workshop on Privacy in the Electronic Society (2006, 2007, 2009)
- ACM Symposium on Applied Computing (2007)

International Award Committee Member:

- SC Awards Europe (2018)
- PET Award (2013)

International Grant Proposal Reviewer:

- EPSRC Peer Review College
- Arts and Humanities Research Council
- Royal Society International Exchanges Panel
- Netherlands Organisation for Scientific Research (NWO)
- Austrian Science Fund (FWF)
- Isaac Newton Institute
- National Science Foundation
- Canada Foundation for Innovation
- Research Council of Norway (Panel Leader for Centre of Excellence Programme)
- European Commission
- United States Air Force Office of Scientific Research (AFOSR)

International Journal Reviewer:

- Proceedings on Privacy Enhancing Technologies (2017, 2018, 2019)
- ACM Transactions on Information & System Security
- IEEE Transactions on Software Engineering
- IEEE/ACM Transactions on Networking
- Identity in the Information Society

Prizes, Awards and Other Honours

2019 Awarded Internet Research Task Force Applied Networking Research Prize
2019 Shortlisted for UCL Provost's Public Engagement Awards
2016 Awarded SIIA CODiE prize for Best Identity & Access Security Solution for DIGIPASS for Apps

2015	Awarded Security Products New Product of the Year for DIGIPASS 760 authentication token
2014	Shortlisted for Cambridge Ring Company of the Year
2011	Shortlisted for the Lloyd's Science of Risk Prize – Chip and PIN is Broken
2010	Awarded the IEEE Award for Outstanding Paper in Security & Privacy – Chip and PIN is Broken
2008	Awarded the IEEE Award for Outstanding Paper in Security & Privacy – Thinking Inside the Box: System-Level Failures of Tamper Proofing
2008	Shortlisted for the Privacy Enhancing Technologies (PET) Award for Outstanding Research – Sampled Traffic Analysis by Internet-Exchange-Level Adversaries
2008	Awarded the European Research Consortium for Informatics and Mathematics (ERCIM) Security and Trust Management Working Group Prize for Best PhD Thesis – Covert Channel Vulnerabilities in Anonymity Systems
2007	Awarded the USENIX Security Prize for Best Student Paper – Keep Your Enemies Close: Distance Bounding Against Smartcard Relay Attacks
2006	Awarded the University of Cambridge Computer Laboratory Prize for the Most Notable Publication of 2006 – Low-Cost Traffic Analysis of Tor
2006	Shortlisted for the Privacy Enhancing Technologies (PET) Award for Outstanding Research – Low-Cost Traffic Analysis of Tor

Invited Talks

37 invited talks given since 2004, inc. 17 at international conferences and 12 keynotes

Jul 18 Royal Society – Privacy Enhancing Technologies

Feb 18 Royal Society UK-Netherlands Bilateral International Meeting – *Transparency enhancing technologies for accountable data science*

Dec 17 European Parliament – Sakharov Debate: is online privacy a human right?

Dec 17 ACI Worldwide – Age of Customer Consent, who owns the Customer Data

Nov 17 University of Edinburgh – Payment Security: Attacks & Defences

Aug 16 IFIP Summer School – Anonymity & Censorship-Free Communication

Nov 15 BT Insights, Adastral Park – *Anonymous Communications*

Sep 15 Information Security Forum – *Privacy by Design*

Dec 14 GCHQ Academic Centers of Excellence conference – *Cyber-security innovation*

Oct 14 [Keynote] Payment Knowledge Forum 2014, London, UK – *Payment Security: Attacks & Defences*

Jun 14 [Keynote] OWASP AppSec Europe 2014, Cambridge, UK – Anonymous Communications and Tor: History and Future Challenges

Nov 13 [Keynote] UK Cyber & Fraud Solutions, British Embassy, Switzerland – *Banking Security*

Sep 13 [Keynote] European Symposium on Research in Computer Security (ESORICS 2013), Royal Holloway, UK – *Security Protocols and the Law: The Case of Chip and PIN*

Sep 13 [Keynote] Quantitative Aspects in Security Assurance, (QASA 2013), Royal Holloway, UK – *Quantifying and Measuring Anonymity*

Mar 13 [Keynote] Open Web Application Security Project (OWASP), Leuven, Belgium – *Banking Security: Attacks and Defences*

Sep 12 [Keynote] Cryptographic Hardware and Embedded Systems (CHES 2012), K.U. Leuven, Belgium – *Banking Security: Attacks and Defences*

Jun 11 [Keynote] Centre for Telematics and Information Technology (CTIT) Symposium, University of Twente – *The Tor Anonymous Communication Network*

Feb 10 [Keynote] Conference on Achieving Sustainable Improvements in the Security of Retail Payments, Federal Reserve Bank of Philadelphia, Philadelphia PA, US – *Chip & PIN: 5 Years On*

Sep 08 [Keynote] Future of Identity in the Information Society (FIDIS)/International Federation for Information Processing (IFIP) Internet Security & Privacy Summer School, Brno, Czech Republic – *The Future of Anonymity and Censor-Free Publishing*

Jun 08 [Keynote] International Workshop on Security and Trust Management, European Research Consortium for Informatics and Mathematics (ERCIM), Trondheim, Norway – *On the Origins of a Thesis*

Sep 07 [Keynote] European Conference on the BSD Family of Operating Systems (EuroBSDCon), Copenhagen, Denmark – *Hot or Not: Fingerprinting Hosts through Clock Skew*

Teaching Career Summary

2015–	Lecturer for University College London <i>COMP0057 Research in Information Security</i> (module co-ordinator, 30 contact hours over 10 lectures: lecturing and development of new lecture material) <i>COMP0064 Dissertation</i> (module co-ordinator, allocation of projects to supervisors and managing assessment; supervising and assessing MSc Information Security projects) <i>COMP0025 Introduction to Cryptography</i> (guest lecturer, on cryptography applied to banking security) <i>COMP0058 Applied Cryptography</i> (guest lecturer, on anonymous communications systems)
2014–	Visiting Lecturer for Computer Science & Engineering at the University of Cambridge <i>Computer Security: Current Applications and Research</i> (lecturing, development of learning material, setting and marking of examinations for MPhil Advanced Computer Science) <i>Security II</i> (Lecturing, development of learning material, setting and marking of examinations for BA Computer Science course in Year 3 of 3)
2002–	<i>Software Engineering</i> (Lecturer, Engineering Department for MEng Engineering in Year 3 of 4) Supervisor for Engineering and Computer Science at the University of Cambridge Tutorials and marking of coursework for: <i>Digital Electronics</i> (BA/MEng Engineering course in Year 1 of 4) <i>Linear Circuits</i> (BA/MEng Engineering course in Year 1 of 4) <i>Dimensional Analysis</i> (BA/MEng Engineering course in Year 1 of 4) <i>Introduction to Security/Security I</i> (BA Computer Science course in Years 2 of 3) <i>Discrete Mathematics</i> (BA Computer Science course in Years 2 of 3) <i>Object Oriented Programming</i> (BA Computer Science course in Years 1 of 3) <i>Security II</i> (BA Computer Science course in Years 3 of 3) Computer Science undergraduate and diploma/MPhil (postgraduate) projects Engineering MEng projects
2014	Visiting Lecturer for Royal Holloway, University of London <i>Smart Cards/Tokens Security and Applications</i> (lecturing and development of new lecture and corresponding formative assessment, on Trusted Execution Environments)
2014	Visiting Lecturer for University College London <i>Web Economics</i> (lecturing and development of new lecture on Online Payments Methods)
2011–2014	Lecturer for Computer Science and Engineering at the University of Cambridge <i>Computer Security: Principles and Foundations</i> (Development of new course, lecturing, development of learning material, setting and marking of examinations for MPhil Advanced Computer Science) <i>Computer Security: Current Applications and Research</i> (Development of new course, lecturing, development of learning material, setting and marking of examinations for MPhil Advanced Computer Science) <i>Security II</i> (Lecturing, development of learning material, setting and marking of examinations for BA Computer Science course in Year 3 of 3) <i>Security I</i> (Lecturing, development of learning material, for BA Computer Science in Year 2 of 3) <i>Software Engineering</i> (Lecturer, Engineering Department for MEng Engineering in Year 3 of 4)

Enterprise/External Engagement

Creator and maintainer of the EMV Lab (<https://emvlab.org/>) research platform in 2009 providing tools for researchers and practitioners in the field of card and payment system security, **currently attracting over 32,000 visits per month.**

Creator of the Tor Browser in 2008 based on my research on Internet privacy. This system is now the flagship product of the Tor Project and the **technology used by the vast majority of Tor's 2 million daily users**.

Creator of the Cronto payment authentication technology, spun out from my banking research in 2007. I served as Chief Security Architect for this company until its acquisition by OneSpan in 2013. The technology is used **banks including market leaders in Germany, Netherlands and Switzerland**.

Short Courses for Professional Development:

- UCL Faculty of Laws – development of a new lecture covering topics including data protection, privacy enhancing technologies, blockchain, and cryptography for legal professionals (2016–)
- SecAppDev Industrial Training course jointly run by K.U. Leuven, Solvay Business School and Trinity College Dublin. Development of new series of lectures: *security economics; anonymity systems requirements and architecture; banking security architecture; ATM and point of sale system security architecture; online banking security* (2010–2014)
- University of Cambridge – development of lectures and practical exercises for an advanced course for information security practitioners in industry: *mobile systems; traffic analysis and anonymity* (2008–2010)

Information Security consultant:

- A hedge fund developing new internal security controls (2018)
- A fund investing in Internet security technologies (2017)
- A start-up developing and commercializing privacy enhancing technologies (2015–)
- A leading academic publisher (2015)
- A company developing secure collaboration tools and services (2013)
- A small company developing secure voice conferencing services (2010)
- Documotion Research, developing secure PIN distribution technologies (2005)
- A leading developer of mobile phone operating systems (2004)

Public and Policy Engagement:

- Member of Advisory Council to the Foundation for Information Policy Research (2019–)
- Contributor to UK Authorised Push Payment Contingent Reimbursement Model consumer protection scheme, including through advising Which? and Age UK on how my research shows how to prevent fraud (2016–2019)
- Author of Home Office standard for secure handling of evidence in the UK justice system (2014–2017)
- Steering group member for Royal Society report on Cybersecurity (2014–2016)
- Technical adviser to House of Commons Science and Technology Committee investigation on the Investigatory Powers Bill (2016)
- Organiser of joint Royal Society & Royal Society of Canada Frontiers of Science meeting on Information and Communication Technologies (2016)
- Editor of Parliamentary Office on Science and Technology report on The Dark Web (2015)
- Witness to the Joint House of Lords and Commons Committee on the Communications Data Bill (2013)
- Editor of Parliamentary Office on Science and Technology report on Digital Identity (2012)
- Frequent meetings with policymakers, working as a Member of the Cambridge Centre for Science and Policy (CSaP) Network (2010–)
- Author of first UK standard on secure distribution of payment card PINs (2005)

Outreach to schools and school-age students:

- Presenter at Royal Institution Engineering Masterclass at University of Cambridge and University of Oxford (2016–)

- Presenter at Royal Institution Computing Masterclass at UCL (2015–)
- Presenter at Royal Institution Maths Masterclass, Cambridge (2011–)
- Supervisor for Nuffield Research Placement (2013)
- Talks to the general public and school students, including during UK National Science & Engineering Week

Work with the media

- Research featured on Computerphile YouTube channel (57,665 views as of October 2019)
- Ran security and source-protection training for Channel 4 investigative journalists
- Author for The New Statesman; article in August 2017
- Author for The Daily Mail; article in August 2016
- Author for The Observer; article in April 2015
- Author for The Conversation; 7 articles with 121,000 readers since February 2015
- Author for The European; article in August 2011
- Reviewer for BBC Tomorrows World online activity on history of science
- Regular interviews with print, online, radio, and TV journalists, including:

The UK – New Scientist, The Times, Guardian, Telegraph, Independent, Financial Times, Daily Mail, BBC News, BBC Watchdog, BBC Newsnight, BBC Fake Britain, ITV Manhunt, Naked Scientists, Rip Off Britain, BBC World Service, The Register, LBC

US - CNN, Wall Street Journal, Huffington Post, Pittsburgh Business Times, ABC News, WIRED

Canada - CBC News

Canada CBS News

Germany - ARD Plusminus, The European, Der Spiegel, Heise

Germany = ARD
Italy = L'Espresso

The Netherlands - VPRO Goudzoekers

The Netherlands – W
Columbia – NTN24

Australia - ABC Background Briefing

Australia - ABC Background
China - Central China TV

Other international outlets – Channel News Asia, Al Jazeera, International Business Times

Selected Publications

For the full list of publications, see <https://murdoch.is/papers>

- Do You See What I See? Differential Treatment of Anonymous Users. Sheharbano Khattak, David Fifield, Sadia Afroz, Mobin Javed, Srikanth Sundaresan, Vern Paxson, Steven J. Murdoch, Damon McCoy. 2016 Network and Distributed System Security Symposium, San Diego, CA, US, 21–24 February 2016.
- Optimising node selection probabilities in multi-hop M/D/1 queuing networks to reduce latency of Tor. Steven Herbert, Steven J. Murdoch, Elena Punskaya. IET Electronics Letters Volume 50, Issue 17, pages 1205–1207, 14 August 2014.
- Impact of Network Topology on Anonymity and Overhead in Low-Latency Anonymity Networks. Claudia Diaz, Steven J. Murdoch, Carmela Troncoso. 10th Privacy Enhancing Technologies Symposium (PETS 2010), Berlin, Germany, 21–23 July 2010.
- An Improved Clock-skew Measurement Technique for Revealing Hidden Services. Sebastian Zander, Steven J. Murdoch. 17th USENIX Security Symposium, San Jose, CA, USA, 28 July–01 August 2008.
- Metrics for Security and Performance in Low-Latency Anonymity Systems. Steven J. Murdoch, Robert N.M. Watson. 8th Privacy Enhancing Technologies Symposium (PETS 2008), Leuven, Belgium, 23–25 July 2008.
- Sampled Traffic Analysis by Internet-Exchange-Level Adversaries. Steven J. Murdoch, Piotr Zieliński. 7th Workshop on Privacy Enhancing Technologies, Ottawa, Canada, 20–22 June 2007.

Exhibit B

No description provided

7h6xs7vpc2qlmm4r2oxfpme3xmj2zvorgu2gwxe6uvq47fk3463qzdad.onion — 6 days, 1 hour ago —

PZA: Quick Search

No description provided

pzaboystoravp2rz.onion – 4 weeks ago –